

AI Visibility Diagnostic

HALVERN INDUSTRIAL • B2B INDUSTRIAL MANUFACTURING • BASELINE DIAGNOSTIC

A worked example of a TKM Studio Answer Share Audit: how three AI engines find, describe and misidentify a mid sized industrial manufacturer, and what causes it.

ABOUT THIS SAMPLE

This is a real client engagement. **Every measurement result in this document is real:** the scores, the counts, the engine behaviour and the technical findings are exactly as recorded. **Every identifying detail has been changed:** company name, country, sector, founding year, facility figures, certification names, named competitors and named references. "Halvern Industrial" and "R. Halvern & Co." are fictional names standing in for the real companies. Published with the client's permission.

SUMMARY

In short: how AI sees Halvern

This report answers one question. When a buyer asks ChatGPT, Gemini or Claude for a manufacturer in this category, does Halvern appear, and if it appears, is it described correctly? We measured it.

The result is three separate problems growing from a single root. ChatGPT does not mention Halvern in most unbranded questions. Gemini mentions it often but merges it with a different company sharing the founder's surname (R. Halvern & Co., a relative's firm), and returns the wrong founding year and the wrong facility data. Claude identifies Halvern correctly, but almost never puts it forward in general questions.

Put simply: when a buyer asks who the best manufacturers in this category are, Halvern is either absent or merged with another company. That is visibility lost before you ever reach the quotation stage.

ENGINE	VISIBILITY (16 UNBRANDED BUYER QUESTIONS)	IDENTITY ACCURACY	MAIN PROBLEM
ChatGPT	Low. About 19 percent, 3 of 16 in both runs.	Neutral, no factual errors.	ABSENCE
Gemini	High. About 50 percent, 8 of 16, stable across both runs.	Corrupted. Merges Halvern with a separate firm.	WRONG IDENTITY
Claude	Near zero. 1 of 16, then 0 of 16.	Correct. Returns accurate facts.	INVISIBILITY

SCORE

Three numbers, read together

The methodology does not produce a single number. Three are read together: how visible you are, how much of that visibility actually works in your favour, and whether you are the right company when you are mentioned.

31/100

ANSWER SHARE
Visibility and favourable mention

26/100

TRUSTED ANSWER SHARE
What remains once identity damage is deducted

85%

IDENTITY INTEGRITY
Share of mentions reaching the correct company

The gap between 31 and 26 is identity damage, and almost all of it comes from Gemini fusing Halvern with the same surname firm. High visibility with low trusted visibility is a worse position than low visibility: the buyer is routed to another company's information under your name.

The most honest number is the score on questions where the brand name never appears. Unbranded Answer Share sits at roughly 12 to 16 out of 100, meaning that when a buyer asks a general category question, Halvern is effectively not there. The score on branded questions looks higher, but the brand is already inside the question, so that figure is structurally inflated and carries little meaning on its own.

On the identity side the sharpest finding is Gemini: roughly one in every four Halvern mentions goes to the wrong company. ChatGPT and Claude get the identity right, but their visibility is very low.

This is a baseline measurement of the current state. Final letter graded scoring is confirmed once certification documents and reference verification arrive. The direction of the findings does not change.

METHOD

What we measured, and how

The measurement was run in a neutral and repeatable way:

- **Three AI engines:** ChatGPT, Gemini and Claude, each on its current model.
- **25 buyer questions:** general category questions, competitor comparisons, real purchasing scenarios and direct brand questions. Most are unbranded, meaning the Halvern name never appears in the question.
- **Full product coverage:** the questions span Halvern's product lines and the different ways a buyer searches, including storage and hangar structures, large clear span builds, load and fire regulation, turnkey and permit exempt solutions, and logistics, agriculture and livestock scenarios. Not one question, but a representative cross section of what a buyer might ask.
- **Two separate clean runs per engine.** Every question asked from zero in a memoryless session, so no previous answer could influence the next.
- **Home market, local language, July 2026.**

The examples shown in this report are illustrative selections from the full set of 25 questions. The complete question list and every raw engine answer are available on request.

FINDINGS

Engine by engine

ChatGPT: absence

ChatGPT mentions Halvern very rarely in unbranded questions (3 of 16, about 19 percent). When it does mention the company it is accurate and neutral, and invents nothing. The problem is not the content, it is the visibility. In a general question Halvern is usually not on the list at all, and the first position typically goes to a large international brand.

Gemini: wrong identity

Gemini mentions Halvern far more often (8 of 16, about 50 percent), but this visibility is deceptive, because the identity is confused. Gemini fuses Halvern Industrial with R. Halvern & Co., a separate company carrying the same founder surname. The consequences:

- It gives the founding year as two years earlier than the real one, sometimes alternating between both years inside a single answer.
- It reports the facilities incorrectly, returning the other company's floor area figure and the wrong location instead of Halvern's real covered and open areas.
- It names the wrong family member as founder.

Worse, the confusion grew between the two runs. A question Gemini separated correctly in the first run was merged in the second. The error is not stable. It is volatile and increasing.

REPUTATION SPILLOVER RISK

This is not only a naming error. The same surname firm has a publicly visible profile on a consumer complaints platform. When AI does not separate the two companies, signals belonging to that firm can be attributed to Halvern. The risk is therefore not visibility alone, it is the mixing of reputations.

Claude: correct identity, but invisible

Claude has the identity right. Asked directly, it returns the correct founding year and the correct founder, states unprompted that R. Halvern & Co. is a separate company, and visits Halvern's official site to verify certifications. But in

unbranded questions Halvern almost never surfaces (1 of 16, then 0 of 16). The information is correct, yet when a buyer asks a general question Halvern does not make the list.

ROOT CAUSE

Why this is happening

Three engines, three different behaviours, one shared root: Halvern's site is missing the structural identity signals AI can trust. What the technical review found:

- **No structured data (JSON-LD / schema) anywhere on the site.** There is no machine readable identity layer for AI to read as "this is an organisation, and this is its name".
- **Core tags missing (meta description, H1).** AI cannot read a clear summary of what the page is.
- **Certifications exist only as images (JPG).** Five international standards are displayed on the site, but because they are pictures they cannot be read as text.
- **Homepage counters (years, projects, square metres) are served as zero.** The real values populate only in the browser, so AI most often sees the raw, empty state.
- **The same surname firm shares the registered address** (a partnership that ended decades ago). With no explicit identity separation, the two merge as if they were one company. This is the root of Gemini's identity confusion.

In short: the engine that crawls broadly (Gemini) finds Halvern but confuses the identities, and the engine that weighs authority and structure (Claude) does not put Halvern on the list at all.

EXAMPLES

Two findings from the set

The reference paradox

Asked which manufacturers in this category work with major industrial groups, Halvern is absent from all three engines. Yet Halvern's own reference page publicly lists exactly such a client, with the project year. The proof is sitting on the site, but AI cannot carry it into the answer to this question. Only Claude found it, and only after visiting the site directly. This is the clearest summary of the visibility problem: the evidence is on the site, it does not reach the AI.

Competitor framing and the unrecognised real competitor

The firms AI puts forward overlap poorly with Halvern's real competitor set. One genuine competitor, a mid sized manufacturer with an active website, is not recognised as a competitor in this category by any engine, and some engines confuse its name with an unrelated software platform. Meanwhile certain budget segment firms rank ahead of Halvern in some answers.

COMPETITIVE PICTURE

Who gets recommended instead

The firms AI puts forward in Halvern's place fall into two groups. Together they show why visibility is an urgent matter rather than a cosmetic one.

Premium and international brands

In category questions the first position generally goes to large international brands. These brands already hold the signals AI trusts: a defined corporate identity, structured and consistent web data, a broad digital footprint. Halvern sits in the same league on work quality and scale, but not on digital visibility. The competition is not happening in the product, it is happening in the identity layer AI can read. Large brands being visible in AI is the signal that this is no longer a preference, it is becoming the category standard.

Budget and price focused local firms

The second group is more uncomfortable. Certain price focused local firms rank ahead of Halvern in some answers. These firms are not in Halvern's quality and scale segment, yet on specific questions they are further forward in AI visibility. The reason is not product superiority, it is which company's data is more accessible to the crawl path AI follows.

THE LESSON

AI visibility is not a quality ranking. With the right optimisation a smaller firm can move ahead of a stronger manufacturer. By the same logic Halvern, with the right identity and data signals, can be visible alongside its premium competitors and stop falling behind the budget segment. This is not a lost quality race. It is an optimisation that has not been done yet.

COMMERCIAL MEANING

What this costs

Buyer research increasingly runs through AI. At that first filtering stage Halvern is currently either invisible or merged with another company. Both lead to the same outcome: the buyer never hears the name and Halvern falls outside the shortlist.

The good news is that this is not a content problem or a reputation problem. It is a visibility and identity problem. Halvern's real data, the long history, the facilities, the certifications, the significant references, is already strong. The only thing missing is that data presented in a form AI can read.

SCOPE OF THE FIX

What can be done

The headings below define the scope of the work TKM Studio carries out in the remediation phase. We do the implementation. The purpose here is to show which areas get addressed. One target: when a buyer asks AI, Halvern comes back with the right identity and the right information.

- **Corporate identity.** The structural work required for AI to recognise Halvern as one clear organisation and stop merging it with R. Halvern & Co.
- **Canonical facts.** Founding year, facilities, certifications and references presented from a single authoritative source in a form AI can read.
- **Certification visibility.** Making the standard information inside the documents machine readable.
- **Crawlability.** Ensuring page and corporate data reach AI intact.
- **Authority.** Surfacing the existing strong references and facts in the right form.

INPUTS

What we need: documents and access

Moving to the remediation phase requires two kinds of input.

Documents

- Certification documents: number, issuing body and validity date for each of the five standards.
- Reference information: publishable ones named, the rest anonymised by region, project type and year.

Access

- **Website and CMS administration.** Schema, meta and content changes are made here.

- **DNS and CDN or WAF panel.** AI crawler access is governed here. At present robots.txt blocks several AI bots while answer engine bots are open. Letting the right bots in and blocking the unnecessary ones requires access at this level.
- **Google Search Console and Bing Webmaster verification,** to accelerate recrawling after the changes and to measure the outcome.

These accesses are needed only for the remediation phase. The diagnostic in this report required none of them. Some access may sit with a third party managing the site, so we clarify together who holds what.

SCOPE AND NEXT STEPS

Why there is no "after" score here

This is a baseline diagnostic, a photograph of the current state. No changes have been made yet, so there is no after measurement, and we will not present one until it exists.

That is deliberate. The sequence is measure, then fix, then measure again, using the same questions, the same engines and the same conditions. An improvement that cannot be demonstrated against a recorded starting point is a claim, not a result. The reason real numbers can appear in this report is that unreal ones never do. The re-measurement window for this engagement is four to eight weeks after the fixes ship.

What can be pointed to now is mechanism and precedent. Each finding above maps to a specific intervention that addresses the mechanism behind it. A machine readable Organization entity is what allows an engine to separate two companies sharing a surname and a registered address. Certifications rendered as text rather than images is what makes a standard quotable inside an answer. Server rendered facts are what an engine reads when it does not execute JavaScript. The technical layer this report prescribes is deployed and live verified on TKM Studio's own properties.

Final letter graded scoring firms up when the certification documents and reference verification arrive. Until then the direction of the findings does not change, only the numeric scores sharpen. The order matters too: the current state is measured first, and changes are made only after the baseline is closed, so improvement can be demonstrated concretely.

AI visibility is not a one time job. Engines and competitors change, so it calls for regular tracking and content. That is a step to discuss after the results of the fix are visible, not now.

We make AI recognize you.
Then we make it work for you.

TKM STUDIO • AI VISIBILITY & ADOPTION STUDIO • MADRID, WORKING WORLDWIDE • TKMSTUDIO.COM
ANONYMISED SAMPLE REPORT • BASELINE DIAGNOSTIC, JULY 2026 • REAL MEASUREMENT, CHANGED IDENTITY